




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : H04J 3/16		A1	(11) International Publication Number: WO 99/09689
			(43) International Publication Date: 25 February 1999 (25.02.99)
(21) International Application Number: PCT/US98/14893 (22) International Filing Date: 17 July 1998 (17.07.98) (30) Priority Data: 60/055,658 14 August 1997 (14.08.97) US 09/061,988 17 April 1998 (17.04.98) US (71) Applicant: MOTOROLA INC. [US/US]; 1303 East Algonquin Road, Schaumburg, IL 60196 (US). (72) Inventors: RUSZCZYK, Chester, A.; Apartment 604, 451 W. Melrose, Chicago, IL 60657 (US). LEE, Whay, Chiou; 16 Michael Way, Cambridge, MA 02141 (US). CHLAMTAC, Imrich; 7575 Frankford Road #325, Dallas, TX 75252 (US). (74) Agents: WOOD, J., Ray et al.; Motorola Inc., Intellectual Property Dept., 1303 East Algonquin Road, Schaumburg, IL 60196 (US).		(81) Designated States: AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GE, GH, GM, HR, HU, ID, IL, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, UZ, VN, YU, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG). Published <i>With international search report. Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i>	

(54) Title: SYSTEM, DEVICE, AND METHOD FOR SCHEDULING IN A COMMUNICATION NETWORK



(57) Abstract

A system, device (214), and method for scheduling in a communication network separates each scheduling job into a committed-rate component (308) and a variable-rate component (310) according to a service class associated with the scheduling job. The committed-rate components are scheduled by creating a map of transmission opportunities that, when repeated at predetermined periodic intervals, provides each committed-rate component with its required number of transmission opportunities within its predetermined delay and jitter constraints. The variable-rate components are scheduled by assigning any unused transmission opportunities in the map to the variable-rate components according to a predetermined scheme.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece			TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	NZ	New Zealand		
CM	Cameroon			PL	Poland		
CN	China	KR	Republic of Korea	PT	Portugal		
CU	Cuba	KZ	Kazakstan	RO	Romania		
CZ	Czech Republic	LC	Saint Lucia	RU	Russian Federation		
DE	Germany	LI	Liechtenstein	SD	Sudan		
DK	Denmark	LK	Sri Lanka	SE	Sweden		
EE	Estonia	LR	Liberia	SG	Singapore		

SYSTEM, DEVICE, AND METHOD FOR SCHEDULING IN A COMMUNICATION NETWORK

Related Applications

- 5 This application claims priority from United States provisional patent application number 60/055,658 entitled System, Device, And Method For Scheduling In A Communication Network filed on August 14, 1997.

Background

10 1. Field of the Invention

The invention relates generally to communication systems, and more particularly to scheduling transmission opportunities in a communication network.

15 2. Discussion of Related Art

In today's information age, there is an increasing need for high-speed communications that provides guaranteed quality of service (QoS) for an ever-increasing number of communications consumers. To that end, communications networks and technologies are evolving to meet current and future demands.

- 20 Specifically, new networks are being deployed which reach a larger number of end users, and protocols are being developed to utilize the added bandwidth of these networks efficiently.

- One technology that has been widely employed and will remain important in the foreseeable future is the shared-medium network. A shared medium
25 network is one in which a single communications channel (the shared channel) is shared by a number of end users such that uncoordinated transmissions from different end users may interfere with one another. Since communications networks typically have a limited number of communications channels, the shared medium network allows many end users to gain access to the network
30 over a single communications channel, thereby allowing the remaining communications channels to be used for other purposes. However, the shared medium network is only feasible when each end user only transmits data intermittently, allowing other end users to transmit during periods of silence.

One problem in a shared medium network involves scheduling end user transmissions to avoid collisions on the shared channel while providing each end user with guaranteed QoS. This scheduling problem becomes extremely complex when a large number of end users is supported. Thus, a need has remained for a scheduling architecture for reducing the complexity of scheduling.

Brief Description of the Drawing

In the Drawing,

FIG. 1 is a block diagram showing an exemplary shared medium network in accordance with a preferred embodiment of the present invention;

FIG. 2 is a block diagram showing a headend unit and an access interface unit in accordance with a preferred embodiment of the present invention;

FIG. 3 is a block diagram showing the Headend Scheduler architecture in accordance with a preferred embodiment of the present invention;

FIG. 4 is a block diagram showing the components of the Headend Scheduler in accordance with a preferred embodiment of the present invention;

FIG. 5A is a diagram showing an exemplary Task Frame map in accordance with an exemplary embodiment of the present invention;

FIG. 5B is a diagram showing an exemplary Task Frame map in accordance with a preferred embodiment of the present invention;

FIG. 6 is a diagram showing an exemplary Task Frame map conceptualized as a number of bins in accordance with a preferred embodiment of the present invention;

FIG. 7 is a block diagram summarizing the relationship between requests, jobs, and tasks in accordance with a preferred embodiment of the present invention;

FIG. 8 is a flow diagram showing logic for scheduling transmission opportunities in accordance with the present invention; and

FIG. 9 is a flow diagram showing logic for scheduling committed-rate components in accordance with the present invention.

Detailed Description

FIG. 1 is a block diagram showing an exemplary shared medium network 100 in accordance with a preferred embodiment of the present invention. The shared medium network 100 allows a number of end users 110, through 110_N to access a remote external network 108 such as the Internet. The shared medium network 100 acts as a conduit for transporting information between the end users 110 and the external network 108.

The shared medium network 100 includes a Headend Unit (HU) 102 that is coupled to the external network 108. The HU 102 is in communication with a plurality of Access Interface Units 104, through 104_N (collectively referred to as "AIUs 104" and individually as an "AIU 104") by means of a shared physical medium 106. Each end user 110 interfaces to the shared medium network 100 by means of an AIU 104. A single AIU 104 may support one or a number of the end users 110.

The shared physical medium 106 includes a plurality of channels over which information can be transferred between the HU 102 and the AIUs 104. In the preferred embodiment, each channel is unidirectional; that is, a particular channel either carries information from the HU 102 to the AIUs 104 or from the AIUs 104 to the HU 102. Those channels that carry information from the HU 102 to the AIUs 104 are typically referred to as "downstream channels." Those channels that carry information from the AIUs 104 to the HU 102 are typically referred to as "upstream channels." In alternative embodiments, these various upstream and downstream channels may, of course, be the same physical channel, for example, through time-division multiplexing/duplexing, or separate physical channels, for example, through frequency-division multiplexing/duplexing.

In the preferred embodiment, the shared medium network 100 is a data-over-cable (DOC) communication system wherein the shared physical medium 106 is a two-way hybrid fiber-optic and coaxial cable (HFC) network. The HU 102 is a headend device typically referred to as a "cable router." The AIUs 104 are cable modems. In other embodiments, the shared physical medium 106 may

be coaxial cable, fiber-optic cable, twisted pair wires, and so on, and may also include air, atmosphere, or space for wireless and satellite communication.

In the shared medium network 100, the downstream channels are situated in a frequency band above approximately 50 MHz. The downstream channels
5 are classified as broadcast channels, since any information transmitted by the HU 102 over a particular downstream channel reaches all of the AIUs 104. Any of the AIUs 104 that are tuned to receive on the particular downstream channel can receive the information.

In the shared medium network 100, the upstream channels are situated in
10 a frequency band between approximately 5 through 42 MHz. Each upstream channel is divided into successive time slots, and is therefore often referred to as a "slotted channel." A slot may be used to carry a protocol data unit including user or control information, or may be further divided into mini-slots used to carry smaller units of information. The upstream channels are classified as shared
15 channels, since only one AIU 104 can successfully transmit in a slot or mini-slot at any given time, and therefore the upstream channels must be shared among the plurality of AIUs 104. If more than one of the AIUs 104 simultaneously transmit on a particular upstream channel, there is a collision that corrupts the information from all of the simultaneously transmitting AIUs 104.

20 In order to allow multiple AIUs 104 to share a single upstream channel, the HU 102 and the AIUs 104 participate in a medium access control (MAC) protocol. The MAC protocol provides a set of rules and procedures for coordinating access by the AIUs 104 to the shared channel. Each AIU 104 participates in the MAC protocol on behalf of its end users. For convenience,
25 each participant in the MAC protocol is referred to as a "MAC User."

A MAC User typically represents a connection supporting a particular end user application. Although it is not elaborated upon, a MAC User could also represent an aggregate of connections that share similar traffic characteristics. In many modern communication networks, each MAC User has specific traffic
30 parameters and QoS requirements, which are agreed upon when the MAC User establishes a connection in the network and which are guaranteed by the network. For example, a MAC User may require a guaranteed minimum amount

of bandwidth, a guaranteed maximum transfer delay for the data it transmits, or a minimum time variation between transmission opportunities provided by the HU 102. In a cell-based network such as an Asynchronous Transfer Mode (ATM) network, a number of parameters are used to characterize the bandwidth

- 5 requirements of each MAC User. Specifically, a set of ATM Traffic Descriptors is used for characterizing the type of traffic generated by the MAC User, and a set of QoS Parameters is used for specifying the network services required by the MAC User.

- 10 The ATM Traffic Descriptors include Peak Cell Rate (PCR), Sustainable Cell Rate (SCR), Minimum Cell Rate (MCR), and Maximum Burst Size (MBS). The PCR is an indication of the maximum data rate (cells per second) that will be generated by the MAC User. The SCR represents the long-term average data rate generated by the MAC User. The MCR is the minimum data rate (cells per second) required by the MAC User. The MBS is the maximum size (number of
- 15 cells) of any burst generated by the MAC User.

- The QoS parameters include Maximum Cell Transfer Delay (MaxCTD), Cell Delay Variation (CDV), and Cell Loss Ratio (CLR). The MaxCTD specifies the maximum delay (including the access delay and propagation delay of the underlying communications network) that will be tolerated by the MAC User. The
- 20 CDV specifies the maximum time variation between data transmission opportunities that will be tolerated by the MAC User. The CLR specifies the MAC User's allowance of cells that may be dropped by the network, in terms of a ratio of lost cells to all the cells transmitted by the MAC User.

- MAC Users having similar traffic characteristics are categorized into
- 25 service categories. ATM defines five service categories, namely continuous bit rate (CBR), real-time variable bit rate (RT-VBR), non-real-time variable bit rate (NRT-VBR), available bit rate (ABR), and unspecified bit rate (UBR). Each MAC User is categorized into one of the service categories.

- CBR connections are characterized by traffic having a fixed bit rate and
- 30 requiring real-time delivery of cells. The ATM Traffic Descriptor for the CBR service category is PCR. The QoS Parameters for the CBR service category are MaxCTD, CDV, and CLR.

RT-VBR connections are characterized by traffic having a variable bit rate and requiring real-time delivery of cells. In a typical RT-VBR connection, the traffic over the connection will have periods where cells are generated in bursts at the PCR and periods of silence where no cells are generated (voice and real-time video are good examples of RT-VBR traffic). The ATM Traffic Descriptors for the RT-VBR service category are PCR, SCR, and MBS. The QoS Parameters for the RT-VBR service category are MaxCTD, CDV, and CLR.

NRT-VBR connections are characterized by traffic having a variable bit rate (similar to RT-VBR connections), but not requiring real-time delivery of cells.

NRT-VBR connections typically require a low CLR.

ABR connections are characterized by traffic requiring a minimum cell rate, but willing to accept additional bandwidth if and when such additional bandwidth becomes available. The ATM Traffic Descriptors for the ABR service category are the MCR and PCR. The MCR is the minimum guaranteed cell rate required by the connection. The PCR is the maximum cell rate at which the connection can transmit if allowed by the network. Consequently, the transmission rate of the ABR connection lies somewhere between the MCR and the PCR. No QoS Parameters are defined for the ABR service category. ABR connections are typically subject to a flow control mechanism with feedback that requires the source of the connection to adapt its rate in response to changing ATM layer transfer characteristics.

UBR connections are not guaranteed any bandwidth and have no QoS requirements. The ATM Traffic Descriptor for the UBR service category is the PCR. The PCR is the maximum cell rate at which the connection can transmit if allowed by the network.

A number of different MAC protocols have been developed for use in the shared medium network 100. These protocols can generally be categorized as contention-free protocols and contention-based protocols. Contention-free protocols, such as time-division multiple-access (TDMA) and round-robin polling, avoid collisions on the shared channel by means of various scheduling methods by authorizing only one MAC User to transmit on an upstream channel at a time. Contention-based protocols, such as certain reservation-based protocols, do not

avoid collisions but instead resolve any collisions that do occur on the shared channel.

In the preferred embodiment, the MAC protocol uses a combination of polling and contention-based reservation for scheduling upstream transmissions.

5 Contention-based reservation requires a MAC User to make a reservation before the HU 102 will allocate bandwidth to the MAC User. The HU 102 provides reservation opportunities in the form of contention mini-slots to the MAC Users by regularly transmitting special control messages (referred to as "entry poll messages") to the AIUs 104 over the downstream channel. The MAC User
10 makes a reservation by transmitting a reservation request message in response to a reservation opportunity provided by the HU 102. Each reservation opportunity typically authorizes multiple MAC Users to transmit in a contention mini-slot, and therefore the MAC Users must contend for a reservation. Not all MAC Users are required to make a reservation. Certain MAC Users are
15 allocated bandwidth regularly without having to make reservations.

FIG. 2 is a block diagram 200 showing the HU 102 and the AIU 104 in greater detail. The block diagram 200 shows the HU 102 in communication with one of the AIUs 104 by means of a downstream channel 230 and an upstream channel 240. The AIU 104 supports at least one End User 110. The HU 102
20 transmits data and control messages to the AIU 104 by means of downstream channel 230, and receives reservation requests and data from the AIU 104 by means of upstream channel 240.

The HU 102 includes a Connection Manager (CM) 215. The CM 215 is responsible for connection admission control for the network. More specifically,
25 the CM 215 is responsible for establishing and terminating connections in the network.

Before an End User such as the End User 110 can communicate over the network, a connection must be established. Therefore, when the End User 110 requests admission to the network, a connection request message is forwarded
30 to the CM 215. The connection request message specifies the bandwidth and QoS requirements for the connection in the form of an ATM Service Category and associated ATM Traffic Descriptors and QoS Parameters. The CM 215

decides whether or not to establish the connection, and thereby admit the End User 110 to the network, based on the ability (or inability) of the network to meet the bandwidth and QoS requirements of the End User 110.

The HU 102 also includes a Headend Scheduler (HES) 214 that is
5 coupled to the CM 215 and to an Adaptive Reservation Manager (ARM) 211.
The HES 214 is responsible for scheduling transmission opportunities for those End Users that have been admitted to the network by the CM 215. When the CM 215 establishes a connection, the CM 215 sends a connection setup message to the HES 214 including a connection identifier and the ATM Service Category
10 and associated ATM Traffic Descriptors and QoS Parameters for the established connection. Similarly, when the CM 215 terminates a connection, the CM 215 sends a connection release message to the HES 214 including a connection identifier for the terminated connection. The HES 214 uses the connection information received from the CM 215 together with feedback information
15 received from the ARM 211 to control the timing of control messages transmitted by the ARM 211.

The ARM 211 is responsible for implementing the MAC protocol in the HU 102. The ARM 211 includes a Reservation Manager (RM) 212 and a Feedback Controller (FC) 213. The RM 212 monitors the contention mini-slots on the
20 upstream channel 240 to determine the result of contention for each contention mini-slot. The RM 212 sends the contention results to the FC 213 and to the HES 214. The FC 213 maintains the state information for each priority class (if multiple priority classes are implemented to support differentiated quality of service), determines the assignment of contention mini-slots for each contention
25 cycle, and formats the control messages to be transmitted on the downstream channel 230. The FC 213 bases the timing of control message transmissions on, among other things, timing information received from the HES 214.

The AIU 104 includes a User Interface (UI) 225 for interfacing with the End User 110. Data transmitted by the End User 110 is received by the UI 225 and
30 stored in a Memory 224. The UI 225 also stores in the Memory 224 a time stamp indicating the arrival time of the data.

The AIU 104 also includes a Control Message Processor (CMP) 222 that is coupled to the Memory 224. The CMP 222 is responsible for processing data and control messages received from the HU 102 by means of Receiver 221. The CMP 222 participates as a MAC User in the MAC protocol on behalf of the End User 110.

FIG. 3 is a block diagram showing the HES 214 architecture in accordance with a preferred embodiment of the present invention. As shown in FIG. 3, the HES 214 includes Request Queues 302. The HES 214 maintains one Request Queue for each ATM service category. When the CM 215 admits a MAC User, the CM 215 sends a request to the HES 214. The request includes a Request Specification and a Service Category. The Request Specification includes an Identifier that identifies the MAC User associated with the request and a Queue Depth that indicates the number of backlogged cells of the MAC User which are awaiting transmission opportunities. The Service Category indicates the ATM service category associated with the MAC User. The request is queued in the Request Queue corresponding to the Service Category.

When the MAC User makes a successful reservation (or if the MAC User is provided transmission opportunities without having to make a reservation), the HES 214 translates the request into a job. The job is essentially a request that requires scheduling by the HES 214. The job includes a Job Specification that includes the Request Specification from the request and also includes the QoS Requirements associated with the MAC User. The job is queued in the Job Buffers 304.

Additionally, an entry polling job is created internally by the HU 102. The entry polling job is not associated with any MAC User. The entry polling job requires transmission opportunities for supporting the contention-based reservation mechanism. For each such transmission opportunity, the HES 102 authorizes one or more MAC Users to transmit a reservation request (rather than data). These reservation opportunities are provided regularly, so each MAC User having data to transmit is able to make a reservation in a timely manner. Like the other jobs, the entry polling job is queued in the Job Buffers 304.

Each job in the Job Buffers 304 is translated into a task or a set of tasks based on the service category, bandwidth, and the QoS requirements indicated in the corresponding Job Specification. Each task corresponds to a unit of data (e.g., an ATM cell) which awaits assignment of a transmission opportunity in the upstream channel. Each task has associated with it a Task Specification that specifies a preferred transmission slot and CDV allowance associated with the task.

The problem of scheduling arises when the HES 214 must allocate transmission opportunities for a number of tasks so that each admitted MAC User can reserve and receive sufficient transmission opportunities to meet its bandwidth, delay, and jitter constraints. When many MAC Users having disparate bandwidth and QoS requirements are supported, the problem of scheduling is extremely complex.

The HES 214 schedules data transmissions by dividing the upstream channel 240 into successive data slots and assigning each data slot to a particular task. Therefore, the HES 214 includes a field in each entry poll message assigning each data slot to a particular task. For convenience, the field is referred to as the Transmission Frame 306.

The Transmission Frame 306 consists of a number of contiguous slots. Each slot in the Transmission Frame 306 corresponds to a data slot on the upstream channel 240. Each slot in the Transmission Frame 306 indicates whether the corresponding data slot is reserved for transmitting MAC User information or for entry polling. If the data slot is reserved for transmitting MAC User information, the slot in the Transmission Frame 306 further indicates which MAC User is permitted to transmit in the data slot.

As shown in FIG. 3, one way of simplifying the HES 214 is to divide the scheduling problem into two sub-problems, one dealing with the scheduling of committed bandwidth components and the other dealing with the scheduling of variable bandwidth components. A committed bandwidth component is one that requires a predetermined amount of transmission opportunities within a given period. In some cases, these transmission opportunities must be provided at regular intervals. CBR traffic is a good example of a committed bandwidth

application, since the HES must provide a steady stream of slots at the PCR within the CDV constraint, which can be met by providing slots at fixed intervals (i.e., regular transmission opportunities). A variable bandwidth component is one that has an elastic demand for transmission opportunities, and does not
5 require a predetermined amount of transmission opportunities within a given period. UBR traffic is a good example of a variable bandwidth application, since the HES 214 does not guarantee the MAC User any bandwidth, and therefore the MAC User can be allocated bandwidth at the convenience of the HES 214 when there is excess bandwidth available.

10 A preferred embodiment exploits the advantages of this architecture by separating each request into a committed-rate component and a variable-rate component. Some requests have only one component (i.e., committed-rate or variable-rate), while others have both components. The committed-rate components are translated into committed-rate jobs 308, while the variable-rate
15 components are translated into variable-rate jobs 310. A committed-rate component scheduler 312 is used to schedule the committed-rate jobs 308, and a variable-rate component scheduler 314 is used to schedule the variable-rate jobs 310. Each scheduler applies a scheduling discipline that is appropriate for its respective component. A Poll Generator 316 places tasks from the committed-
20 rate component scheduler 312 and the variable-rate component scheduler 314 into Transmission Frames 306. The committed-rate jobs are guaranteed slots within the Transmission Frames 306. The variable-rate jobs are provided any extra slots within the Transmission Frames 306 that are not assigned to committed-rate jobs.

25 The CBR service category includes only a committed-rate component. The CBR MAC User is continuously allocated regular transmission opportunities at the PCR until the CBR connection is terminated. The transmission opportunities are allocated within the MAC User's delay and jitter constraints.

The RT-VBR service category includes only a committed-rate component.
30 A number of alternative embodiments for scheduling RT-VBR traffic are described in United States provisional patent application number 60/055,658 entitled System, Device, And Method For Scheduling In A Communication

Network filed on August 14, 1997, incorporated herein by reference in its entirety. One embodiment, described starting at Page 11, Line 16, treats certain RT-VBR connections as if they were CBR connections. This may be appropriate for certain RT-VBR connections, specifically those with a ratio of SCR to PCR exceeding a predetermined threshold (referred to as "CBR-like" connections), but may be inefficient for other RT-VBR connections (referred to as "bursty" connections). Two approaches are described for scheduling "bursty" RT-VBR traffic. In a first embodiment, described starting at Page 12, Line 15, the RT-VBR MAC User is allocated regular transmission opportunities at the PCR while the MAC User is active and at a zero rate while the MAC User is inactive. In a second embodiment, described starting at Page 14, Line 13, the RT-VBR MAC User is allocated regular transmission opportunities at the PCR while the MAC User is active and at a lower rate (preferably the SCR) while the MAC User is inactive. It is worth noting that, even though the bandwidth requirements for a "bursty" RT-VBR connection can vary, the "bursty" RT-VBR connection does not include a variable-rate component because the RT-VBR MAC User still requires a predetermined amount of transmission opportunities within any given period. The transmission opportunities are allocated within the MAC User's delay and jitter constraints.

The NRT-VBR service category includes both a committed-rate component and a variable-rate component. The NRT-VBR MAC User is allocated transmission opportunities at the SCR (i.e., the committed-rate component) and is allocated additional transmission opportunities when the amount of data queued for transmission by the MAC User exceeds a predetermined threshold (i.e., the variable-rate component). The signaling of queue depth information by the MAC User to the HES 214 can be done, for example, by including status information along with data transmissions (often referred to as "piggy-backing"). For NRT-VBR, there are no specific delay and jitter constraints for the committed-rate component.

The ABR service category includes both a committed-rate component and a variable-rate component. The ABR user is allocated transmission opportunities at the MCR (i.e., the committed-rate component) and is allocated additional

transmission opportunities if there are any transmission opportunities available after the assignments for the committed-rate jobs (i.e., the variable-rate component). For ABR, there are no specific delay and jitter constraints for the committed-rate component.

- 5 The UBR service category includes only a variable-rate component. The UBR user is allocated "leftover" transmission opportunities if there are any.

Like the user requests, the entry polling job is separated into a committed-rate component and a variable-rate component. In the preferred embodiment, the entry polling job includes both a committed-rate component for providing
10 regular reservation opportunities and a variable-rate component for providing additional reservation opportunities. These additional reservation opportunities are optional, although they can improve MAC protocol performance during collision resolution. In an alternative embodiment, the entry polling job includes only a committed-rate component.

- 15 One technique for scheduling for multi-class users with different bandwidth and QoS requirements is taught by U.S. Patent Number 5,528,513 entitled "Scheduling and Admission Control Policy for a Continuous Media Server" issued June 18, 1996 in the name of Vaitzblit et al. (hereinafter Vaitzblit). In Vaitzblit, which applies to a continuous media file server, three classes of
20 users are supported: general-purpose, real-time, and isochronous. The general-purpose class supports preemptive tasks that are suitable for low-priority background processing. This class is, however, granted a minimum CPU processing quantum. The real-time class represents users that require guaranteed throughput and bounded delay. This class is not preemptive, but a
25 preemption window is provided in which higher priority isochronous users may be scheduled. The isochronous class represents users that require periodic transmissions, performance guarantees for throughput, bounded latency, and low jitter.

- 30 In Vaitzblit, isochronous tasks are assigned the highest priority and are scheduled first followed by real-time and general-purpose tasks. Isochronous tasks run periodically and are invoked by a timer interrupt set for each task. After the isochronous tasks have been scheduled, the scheduler alternates between

the real-time tasks and the general-purpose tasks using a weighted round-robin scheme.

5 Isochronous tasks associated with users requiring different rates are further prioritized in a rate-monotonic manner; that is, a user with a higher rate has a higher priority. These tasks run periodically and are invoked by timer interrupts set for their corresponding time periods. The scheduler executes isochronous tasks from a ready queue, in which the isochronous tasks are arranged in the order of decreasing priority. An isochronous task is inserted in its appropriate place on the ready queue upon arrival at the server. Isochronous
10 tasks arrive at the expiration of every periodic timer. A unique periodic timer exists in the system for each distinct period among all the admitted isochronous tasks. All isochronous tasks with the same period are executed when the timer associated with the period expires.

When an isochronous task arrives at the server, the scheduler determines
15 if a currently running task needs to be preempted. If the currently running task has a lower priority than that of the arriving task, it is preempted at the next preemption window by the incoming task from the head of the queue. The preempted task is queued to the head of the queue associated with the task, since by virtue of its being a previously running committed-rate task, it had to
20 have the highest priority of all the tasks waiting on the queue.

The Vaitzblit approach for scheduling tasks, which correspond to data units pending transmission, suffers from its real-time, preemptive nature. Specifically, with preemption allowed, it is difficult to support bandwidth and QoS guarantees. In addition, timer management for scheduling isochronous tasks
25 can be cumbersome if there is a large number of distinct rates, since each rate uses a separate timer.

A preferred embodiment of the present invention uses a non-preemptive scheduling technique for scheduling tasks. Specifically, the committed-rate component scheduler 312 creates a map of transmission opportunities that,
30 when repeated at periodic intervals, provides each committed-rate job with its required bandwidth within any delay and jitter constraints. The map consists of a number of slots that are initialized based on the bandwidth and QoS

requirements of the committed-rate jobs. Once created for a given set of jobs, the map typically does not change until either a job is added, a job is deleted, or the bandwidth and/or QoS requirements for an existing job changes. By creating a map of transmission opportunities, the committed-rate jobs can be scheduled *a priori*, not in real-time on a frame-by-frame basis.

Typically, not all slots in the map will be used for scheduling committed-rate jobs. Those slots that are not used for scheduling committed-rate jobs are available for allocation to variable-rate jobs. The map of transmission opportunities has, by definition, a repeatable pattern of slot assignments. It is desirable to use the smallest repeatable pattern to avoid long transmission cycles. When the smallest repeatable pattern is too large to be carried in a single entry poll message, the map may be divided into smaller equal portions (referred to as Task Frames). The Poll Generator includes one Task Frame in each Transmission Frame 306.

FIG. 4 is a block diagram showing the components of the HES 214 in accordance with a preferred embodiment of the present invention. As shown in FIG. 4, the committed-rate component scheduler 312 includes a Mapper 402 and a number of Task Frames 404. The Mapper 402 translates the committed-rate jobs 308 into tasks and places the tasks into the Task Frames 404 in a manner described in detail below.

The variable-rate component scheduler 314 includes a Prioritizer 406 and a Task Queue 408. The Prioritizer 406 translates the variable-rate jobs 310 into tasks and places the tasks into the Task Queue 408. The Prioritizer 406 preferably uses a weighted round-robin scheduling discipline for prioritizing the variable-rate jobs 310. The weights used for the weighted round-robin scheduling can be a function of ATM service categories, queue depths, and other relevant characteristics of the variable-rate jobs 310, and are subject to change over time.

The Poll Generator 316 includes one Task Frame 404 in each Transmission Frame 306. The Poll Generator 316 selects a Task Frame 404 from among the number of Task Frames 404 on a strict round-robin basis. The Poll Generator 316 then fills in any empty slots in the Transmission Frame 306

with variable-rate tasks from the Task Queue 408, preferably selected on a first-come-first-served basis.

A key element of the present invention is the scheduling of the committed-rate jobs 308 into the Task Frames 404. The number of slots that comprise the map is generally a function of the set of distinct bandwidth requirements supported by the system. When the job having the lowest bandwidth requirement has a period that is an integral multiple of each other distinct period supported by the system, the lowest bandwidth job will have the longest period for a repeating pattern of slots within the map, and all other jobs will have repeatable patterns within the same period. In one embodiment, the number of Task Frames 404 that collectively constitute the repeatable pattern is determined by the job having the highest bandwidth requirement, and is preferably chosen so that every committed-rate job 308 is allocated no more than one slot per Task Frame 404.

The rate R associated with each committed-rate job 308, in terms of cells per second, can be translated into a transmission period S in terms of a number of transmission slots (i.e., one cell must be transmitted every S slots, where S may not be an integer). Supposing that t is the time it takes, in seconds, to transmit a cell in a transmission slot, then the period in seconds would be $S \cdot t$. In order to meet the rate R precisely, one cell must be transmitted each period. Therefore, $R \cdot S \cdot t = 1$. It follows that the period S is derived from $1/(R \cdot t)$. If $1/(R \cdot t)$ is an integer, then each periodic cell transmission falls directly on a slot boundary. However, if $1/(R \cdot t)$ is not an integer, then cell transmissions cannot be scheduled with a period of $1/(R \cdot t)$, since each periodic cell transmission would not fall directly on a slot boundary.

One alternative is to set S equal to the "floor" of $1/(R \cdot t)$ (i.e., the largest integer that is smaller than $1/(R \cdot t)$). The consequence of this is that the transmission opportunities are scheduled more frequently than necessary to support the rate R . The implication of this rounding error is that certain cells do not carry a full payload of data and therefore have to be padded appropriately with "idle bits".

A second alternative is to set S equal to the "ceiling" of $1/(R \cdot t)$ (i.e., the smallest integer that is larger than $1/(R \cdot t)$). The consequence of this is that the transmission opportunities are scheduled less frequently than required to support the rate R . The implication of this rounding error is that certain cells may be dropped if they cannot be buffered or transmitted within their delay and jitter constraints. A preferred embodiment selects S to be the "floor" of $1/(R \cdot t)$ if $1/(R \cdot t)$ is not an integer.

In accordance with the present invention, if all the committed-rate jobs have a common transmission rate R , a map of size S as determined above may be used for scheduling. In this case, each user requires a single slot in the map (i.e., each user requires one slot per period). The position of the slot assigned to a user may depend on its delay and jitter constraints. Since all of the users share the same periodicity, only one Task Frame of size S is required.

The above non-preemptive scheduling technique can also be used to support users having different transmission rates. Each distinct rate associated with a committed-rate job can be translated into a period in a manner as described above. Let the set of distinct rates associated with the committed-rate jobs be $\{R_1, R_2, \dots, R_m\}$, and the corresponding distinct periods be $\{S_1, S_2, \dots, S_m\}$. Now, let S^* be defined to be the Least Common Multiple (LCM) among S_i , for $i = 1, 2, \dots, m$. If a Task Frame of size S^* is used, it is guaranteed that the transmission pattern within the Task Frame is a repeatable pattern, and also the smallest repeatable pattern.

If every distinct period is an integral divisor of the largest period, S_{\max} , then S_{\max} is simply the LCM among the set of distinct periods, and a Task Frame of size S_{\max} may be used. In general, S^* , the LCM among a set of distinct periods, may be considerably larger than S_{\max} . If a Task Frame of size S^* is used, it may be impossible to meet the delay constraints of some users. One preferred embodiment is to use frames of size S_{\min} , which is the smallest of S_i , for $i = 1, 2, \dots, m$, such that there is still at least one assignment in each Task Frame. In this way, no user is unnecessarily delayed. When smaller frames are used, the number of Task Frames in the map is $K = S^*/S_{\min}$.

The above scheduling technique can be demonstrated by example. In the example, there are three jobs to be scheduled. Job number one has period 4. Job number two has period 8. Job number three has period 16. Using the above notation, the scheduling problem can be modeled as $S_{\max} = 16$, $S_{\min} = 4$, and $S^* = 16$ (i.e., the LCM of 4, 8, and 16).

In one embodiment, shown in FIG. 5A, a single Task Frame of size $S^* = 16$ is used. This guarantees that the transmission pattern within the Task Frame is a repeatable pattern, and also the smallest repeatable pattern. In this case, job number one is allocated four slots within the Task Frame, job number two is allocated two slots within the Task Frame, and job number three is allocated one slot within the Task Frame.

In a preferred embodiment, shown in FIG. 5B, Task Frames of size $S_{\min} = 4$ are used. Because the smallest repeating pattern is $S^* = 16$ slots, $K = 4$ Task Frames are needed. In this case, job number one is allocated one slot in each Task Frame, job number two is allocated one slot in every second Task Frame, and job number three is allocated one slot in every fourth Task Frame.

When scheduling the committed-rate jobs 308 into the Task Frames 404, it is more difficult to schedule those jobs that have delay and jitter constraints than to schedule those jobs that do not have delay and jitter constraints. This is because the delay and jitter constraints dictate which slots in the map can and cannot be used for a particular task (i.e., which slots fall within the delay and jitter constraints and which slots do not). Furthermore, when scheduling the committed-rate jobs 308 into the Task Frames 404, it is generally more difficult to schedule high-bandwidth jobs than to schedule low-bandwidth jobs.

Therefore, as shown in FIG. 4, a preferred embodiment separates the committed-rate jobs 308 into QOS Jobs 410 (i.e., jobs having delay and jitter constraints) and NO QOS Jobs 412 (i.e., jobs having no delay and jitter constraints). The QOS Jobs 410 are prioritized according to their data rate requirements, where the highest rate job is given the highest priority. Likewise, the NO QOS Jobs 412 are prioritized according to their data rate requirements, where the highest rate job is given the highest priority. Such a rate-monotonic

queuing discipline is similar to the rate-monotonic queuing discipline for the ready queue in Vaitzblit's approach for scheduling isochronous tasks.

The committed-rate component scheduler 312 first schedules the QOS Jobs 410, beginning with the highest priority job, in such a way that each of the jobs is allocated its required bandwidth within its required delay and jitter constraints. For these jobs, the problem of allocating slots in the Task Frames 404 can be modeled as a kind of bin packing problem. As shown in FIG. 6, the plurality of Task Frames 404 can be conceptualized as a number of bins. If jitter is not permitted, then each "column" consisting of one slot from each of the Task Frames 404 is the equivalent of a bin. If jitter is permitted, then a plurality of adjacent "columns" may be the equivalent of a bin.

A one-dimensional packing algorithm is discussed in D.S. Johnson, A. Demers, J.D. Ullman, M.R. Garey, and R.L. Graham, Worst-Case Performance Bounds For Simple One-Dimensional Packing Algorithms, Society for Industrial and Applied Mathematics Journal of Computing, Vol. 3, No. 4, December 1974. A more complex bin packing problem for supporting dynamic packing is discussed in E.G. Coffman, Jr., M.R. Carey, and D.S. Johnson, Dynamic Bin Packing, Society for Industrial and Applied Mathematics Journal of Computing, Vol. 12, No. 2, May 1983. In short, the bin packing problem is typically one of packing a number of different size items into a fixed space. In general, a satisfactory result is obtained by first packing the largest items, and then filling in with smaller and smaller items. Improved results (albeit at the expense of time and complexity) can be obtained by allowing items to be removed, rearranged, and re-packed.

The bin packing problem faced by the committed-rate component scheduler 312 has parallels to the dynamic bin packing problem, requiring the "packing" of a number of jobs having different bandwidth requirements into a fixed number of slots. The committed-rate component scheduler 312 typically begins packing the bins to meet the bandwidth and QoS requirements of the highest priority job, and fills in with the remaining jobs in the order of decreasing bandwidth requirements. When allocating slots within the Task Frames 404, it is preferable that the slots be distributed evenly across each column of the Task

Frames 404. If at any time the committed-rate component scheduler 312 cannot meet the delay and jitter requirements of a job using the slots that are available, the committed-rate component scheduler 312 modifies the allocation of slots. For example, the committed-rate component scheduler 312 may move an
5 existing allocation from one slot to an acceptable neighboring slot (i.e., a slot within the delay and jitter constraints of the corresponding job) in order to free up a slot that is within the jitter constraints of the job being scheduled. In a worst-case scenario, the committed-rate component scheduler 312 must unpack, rearrange, and re-pack some or all of the jobs. The detailed formulation of the
10 above packing problem and heuristic solutions are beyond the scope of the present invention.

Once the QOS Jobs 410 have been scheduled, the committed-rate component scheduler 312 schedules the NO QOS Jobs 412, beginning with the highest priority job, in such a way that each of the jobs is allocated its required
15 bandwidth (delay and jitter is not an issue). Although the distribution of slots for the NO QOS Jobs 412 is less critical, it is still preferable that the NO QOS Jobs 412 be evenly distributed throughout the Task Frames 404. This provides a fairly constant number of empty slots in each Transmission Frame 306 that are available for the variable-rate jobs 310.

20 Each Transmission Frame 306 generated by the Poll Generator 316 is transmitted to the MAC Users in an entry poll message. MAC Users having data to transmit are permitted to transmit in designated data slots. The HU 102 monitors the data slots to determine a MAC User status.

A MAC User that transmits data when provided a transmission opportunity
25 is typically considered to be "active." A MAC User that transmits no data when provided a transmission opportunity is typically considered to be "inactive" (unless the MAC User is not required to make a reservation, in which case the MAC User remains "active"). However, additional MAC User status information may be available to the HU 102. For example, a MAC User that transmits data
30 may include an indication whether or not it has more data to transmit (often referred to as "piggy-backing"). An indication that the MAC User has no more

data to transmit may be used to indicate an "inactive" status, even though the MAC User transmitted data.

The HES 214 typically continues providing transmission opportunities to "active" MAC Users. The HES 214 typically stops providing transmission opportunities to "inactive" MAC Users. In this latter case, the HES 214 considers the associated job to be "fulfilled," and prunes (deletes) the associated job from the Job Buffers 304. A job associated with a MAC User is considered "fulfilled" if certain idle conditions are met. The idle conditions are defined with respect to the ATM service category and QoS requirements for the job.

For CBR, NRT-VBR, and ABR jobs, the job is considered "fulfilled" if and only if the corresponding connection is terminated. Otherwise, the job (or jobs) remains in the Job Buffers 304, and the MAC User continues to be allocated bandwidth.

RT-VBR jobs may or may not require pruning, depending on the technique used for supporting RT-VBR jobs as described above. In one embodiment, RT-VBR jobs are allocated bandwidth at the PCR while "active" and at zero while "inactive." Such RT-VBR jobs may be considered "active" when the MAC User makes a successful reservation, and "inactive" when the MAC User fails to transmit data in response to a predetermined number of consecutive transmission opportunities (e.g., two). Furthermore, such RT-VBR jobs may be considered "fulfilled" when the MAC User is "inactive." In this embodiment, the RT-VBR job is dynamically added to the Job Buffers 304 when the MAC User becomes "active," and pruned from the Job Buffers 304 when the MAC user becomes "inactive."

In an alternative embodiment, RT-VBR jobs are allocated bandwidth at the PCR while "active" and at a lower (non-zero) rate while "inactive." Such RT-VBR jobs may be considered "active" at all times (as long as the connection exists). In this embodiment, the RT-VBR job remains in the Job Buffers 304 indefinitely, although the relative priority of such a RT-VBR job (based on bandwidth requirements as described above) may change each time the rate changes. Such rate changes prompt the committed-rate component scheduler 312 to update the Task Frames 404 based on the new rate information.

For UBR users, a job is considered "fulfilled" when it has no backlogged data.

Entry polling jobs must remain active at all times, and are therefore never pruned.

5 When a job is pruned from the Job Buffers 304, it is translated back to a request and queued in the appropriate Request Queue 302. At the same time, all tasks corresponding to the job are pruned from the Task Frames 404 and/or the Task Queue 408. If the request itself is "fulfilled," for example, when the connection is terminated, then the request is pruned from the Request Queues
10 302.

FIG. 7 is a block diagram summarizing the relationship between requests, jobs, and tasks in accordance with a preferred embodiment of the present invention. New requests 702 from the CM 215 are queued in Request Queues 302. The requests 702 are translated into new jobs 704 and stored in Job
15 Buffers 304. Also, an entry polling job 706 is stored in Job Buffers 304. The jobs 704 and 706 are translated into tasks 708, which are then allocated transmission opportunities. Fulfilled jobs 710 are pruned from the Job Buffers 304 and translated back into requests that are queued in Request Queues 302. Fulfilled requests 712 are pruned from the Request Queues 302.

20 FIG. 8 is a flow diagram showing logic for scheduling transmission opportunities in accordance with the present invention. The logic begins in step 802, and proceeds to separate each of a plurality of scheduling jobs into a committed-rate component and a variable-rate component, in step 804. The committed-rate component requires a predetermined number of transmission
25 opportunities within a given amount of time, while the variable-rate component requires an undetermined number of transmission opportunities within the given amount of time. The logic schedules the committed-rate components by creating a map of transmission opportunities that, when repeated at predetermined periodic intervals, provides each committed-rate component with its required
30 number of transmission opportunities within its predetermined delay and jitter constraints, in step 806. The logic schedules the variable-rate components by assigning any unused transmission opportunities in the map to the variable-rate

components according to a predetermined scheme, in step 808. The logic terminates in step 899.

FIG. 9 is a flow diagram showing logic for scheduling committed-rate components in accordance with the present invention. The logic begins in step 902, and proceeds to separate the committed-rate components into committed-rate components requiring delay and jitter constraints and committed-rate components requiring no delay and jitter constraints, in step 904. The logic prioritizes the committed-rate components requiring delay and jitter constraints based on bandwidth requirements in step 906, and prioritizes the committed-rate components requiring no delay and jitter constraints based on bandwidth requirements in step 908. The logic then maps the committed-rate components requiring delay and jitter constraints into the map in order of priority in step 910, and subsequently maps the committed-rate components requiring no delay and jitter constraints into the map in order of priority in step 912. The logic terminates in step 999.

All logic described herein can be embodied using discrete components, integrated circuitry, programmable logic used in conjunction with a programmable logic device such as a Field Programmable Gate Array (FPGA) or microprocessor, or any other means including any combination thereof.

Programmable logic can be fixed temporarily or permanently in a tangible medium such as a read-only memory chip, a computer memory, a disk, or other storage medium. Programmable logic can also be fixed in a computer data signal embodied in a carrier wave, allowing the programmable logic to be transmitted over an interface such as a computer bus or communication network.

All such embodiments are intended to fall within the scope of the present invention.

The present invention may be embodied in other specific forms without departing from the essence or essential characteristics. The described embodiments are to be considered in all respects only as illustrative and not restrictive.

We claim:

1. A method for scheduling transmission opportunities in a communication network, the method comprising the steps of:
 - 5 separating each of a plurality of scheduling jobs into a committed-rate component and a variable-rate component, the committed-rate component requiring a predetermined number of transmission opportunities within a given amount of time, the variable-rate component requiring an undetermined number of transmission opportunities within the given amount of time;
 - 10 scheduling the committed-rate components by creating a map of transmission opportunities that, when repeated at predetermined periodic intervals, provides each committed-rate component with its required number of transmission opportunities within its predetermined delay and jitter constraints; and
 - 15 scheduling the variable-rate components by assigning any unused transmission opportunities in the map to the variable-rate components according to a predetermined scheme.
2. The method of claim 1 wherein the step of scheduling the committed-rate components comprises:
 - separating the committed-rate components into committed-rate components requiring delay and jitter constraints and committed-rate components requiring no delay and jitter constraints;
 - prioritizing the committed-rate components requiring delay and jitter constraints based on bandwidth requirements;
 - 25 prioritizing the committed-rate components requiring no delay and jitter constraints based on bandwidth requirements;
 - mapping the committed-rate components requiring delay and jitter constraints into the map in order of priority; and
 - 30 subsequently mapping the committed-rate components requiring no delay and jitter constraints into the map in order of priority; and wherein the step of scheduling the variable-rate components comprises:

prioritizing the variable-rate components using a weighted round-robin scheduling discipline, wherein the weights used for the weighted round-robin scheduling are a function of ATM service categories, queue depths, and other relevant characteristics of the variable-rate components, and are subject to
5 change over time; and

storing the prioritized variable-rate components in a queue.

3. The method of claim 2 wherein the steps of separating and scheduling comprise:

10 separating each continuous bit rate scheduling job into only a committed-rate component requiring delay and jitter constraints, and scheduling transmission opportunities for the committed-rate component of each continuous bit rate scheduling job at its predetermined peak cell rate within its predetermined delay and jitter constraints;

15 separating each real-time variable bit rate scheduling job into only a committed-rate component requiring delay and jitter constraints, and one of:
scheduling transmission opportunities for the committed-rate component of each real-time variable bit rate scheduling job at its predetermined peak cell rate within its predetermined delay and jitter constraints;

20 scheduling transmission opportunities for the committed-rate component of each real-time variable bit rate scheduling job at its predetermined peak cell rate within its predetermined delay and jitter constraints when the committed-rate component is active, and at a zero rate when the committed-rate component is inactive; and

25 scheduling transmission opportunities for the committed-rate component of each real-time variable bit rate scheduling job at its predetermined peak cell rate within its predetermined delay and jitter constraints when the committed-rate component is active, and at a lower (non-zero) rate within its predetermined delay and jitter constraints when the committed-rate component is
30 inactive;

separating each non-real-time variable bit rate scheduling job into a committed-rate component requiring no delay and jitter constraints and a

variable-rate component, and scheduling transmission opportunities for the committed-rate component of each non-real-time variable bit rate scheduling job at its predetermined sustainable cell rate and scheduling transmission opportunities for the variable-rate component of each non-real-time variable bit rate scheduling job whenever additional transmission opportunities are needed above its predetermined sustainable cell rate based on its backlog of data to be transmitted;

separating each available bit rate scheduling job into a committed-rate component requiring no delay and jitter constraints and a variable-rate component, and scheduling transmission opportunities for the committed-rate component of each available bit rate scheduling job at its predetermined minimum cell rate and scheduling transmission opportunities for the variable-rate component of each available bit rate scheduling job up to its predetermined peak cell rate;

separating each unspecified bit rate scheduling job into only a variable-rate component and scheduling transmission opportunities for the variable-rate component of each unspecified bit rate scheduling job whenever transmission opportunities are available; and

separating each entry polling scheduling job into both a committed-rate component and a variable-rate component and scheduling transmission opportunities for the committed-rate component of the entry polling scheduling job at regular intervals and scheduling transmission opportunities for the variable-rate component of the entry polling scheduling job whenever transmission opportunities are available.

4. The method of claim 1 wherein each committed-rate component has a transmission period S and wherein the map comprises a predetermined number of slots at least equal to a least common multiple S^* of the transmission periods S .

5. The method of claim 1 further comprising the step of generating an entry poll message including an assignment of transmission opportunities, wherein the

map includes a number of task frames and wherein the step of generating the entry poll message comprises:

- including one of the number of task frames in the entry poll message; and
 - assigning any unused transmission opportunities in the task frame to the
- 5 variable-rate components according to a predetermined scheme.

6. An apparatus comprising a computer usable medium having computer readable program code means embodied therein for scheduling transmission opportunities in a communication network, the computer readable program code
- 10 means comprising:

- computer readable program code means for separating each of a plurality of scheduling jobs into a committed-rate component and a variable-rate component, the committed-rate component requiring a predetermined number of transmission opportunities within a given amount of time, the variable-rate
- 15 component requiring an undetermined number of transmission opportunities within the given amount of time;

- computer readable program code means for scheduling the committed-rate components by creating a map of transmission opportunities that, when repeated at predetermined periodic intervals, provides each committed-rate
- 20 component with its required number of transmission opportunities within its predetermined delay and jitter constraints; and

- computer readable program code means for scheduling the variable-rate components by assigning any unused transmission opportunities in the map to the variable-rate components according to a predetermined scheme.
- 25

7. A computer data signal embodied in a carrier wave, wherein embodied in the computer data signal are computer readable program code means for scheduling transmission opportunities in a communication network, the computer readable program code means comprising:

- 30 computer readable program code means for separating each of a plurality of scheduling jobs into a committed-rate component and a variable-rate component, the committed-rate component requiring a predetermined number of

transmission opportunities within a given amount of time, the variable-rate component requiring an undetermined number of transmission opportunities within the given amount of time;

5 computer readable program code means for scheduling the committed-rate components by creating a map of transmission opportunities that, when repeated at predetermined periodic intervals, provides each committed-rate component with its required number of transmission opportunities within its predetermined delay and jitter constraints; and

10 computer readable program code means for scheduling the variable-rate components by assigning any unused transmission opportunities in the map to the variable-rate components according to a predetermined scheme.

8. A device for scheduling transmission opportunities in a communication network, the device comprising:

15 logic for separating each of a plurality of scheduling jobs into a committed-rate component and a variable-rate component, the committed-rate component requiring a predetermined number of transmission opportunities within a given amount of time, the variable-rate component requiring an undetermined number of transmission opportunities within the given amount of time;

20 logic for scheduling the committed-rate components by creating a map of transmission opportunities that, when repeated at predetermined periodic intervals, provides each committed-rate component with its required number of transmission opportunities within its predetermined delay and jitter constraints; and

25 logic for scheduling the variable-rate components by assigning any unused transmission opportunities in the map to the variable-rate components according to a predetermined scheme.

9. A scheduler for scheduling transmission opportunities in a communication network, the scheduler comprising:

30 separating logic including logic for separating each of a plurality of scheduling jobs into a committed-rate job and a variable-rate job, the committed-

rate job requiring a predetermined number of transmission opportunities within a given amount of time, the variable-rate job requiring an undetermined number of transmission opportunities within the given amount of time;

a committed-rate component scheduler, operably coupled to the separating logic, the committed-rate component scheduler comprising a mapper including logic for creating a map of transmission opportunities that, when repeated at predetermined periodic intervals, provides each committed-rate job with its required number of transmission opportunities within its predetermined delay and jitter constraints; and

a variable-rate component scheduler, operably coupled to the separating logic, the variable-rate component scheduler comprising a prioritizer including logic for prioritizing the variable-rate jobs.

10. A system comprising:


a headend unit in communication with an access interface unit by means of a shared medium, the headend unit having a scheduler for scheduling transmission opportunities, the scheduler comprising:

separating logic including logic for separating each of a plurality of scheduling jobs into a committed-rate job and a variable-rate job, the committed-rate job requiring a predetermined number of transmission opportunities within a given amount of time, the variable-rate job requiring an undetermined number of transmission opportunities within the given amount of time;

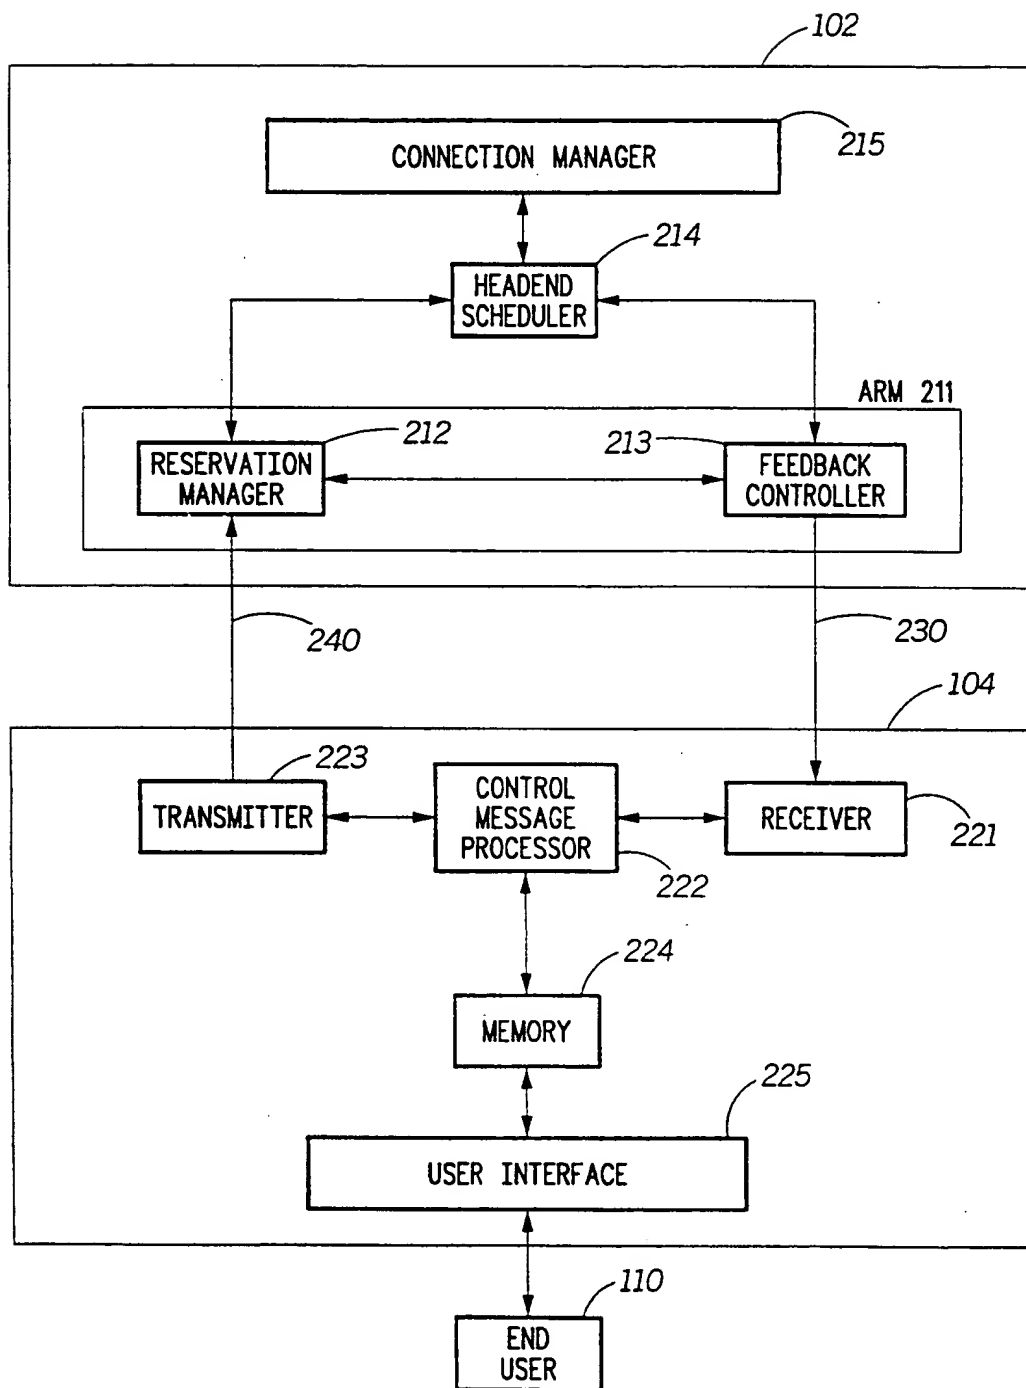
a committed-rate component scheduler, operably coupled to the separating logic, the committed-rate component scheduler comprising a mapper including logic for creating a map of transmission opportunities that, when repeated at predetermined periodic intervals, provides each committed-rate job with its required number of transmission opportunities within its predetermined delay and jitter constraints; and

a variable-rate component scheduler, operably coupled to the separating logic, the variable-rate component scheduler comprising a prioritizer including logic for prioritizing the variable-rate jobs; and

1/7

*FIG. 1*

2/7



200
FIG. 2

3/7




FIG. 3

4/7




FIG. 4




FIG. 5A




FIG. 5B


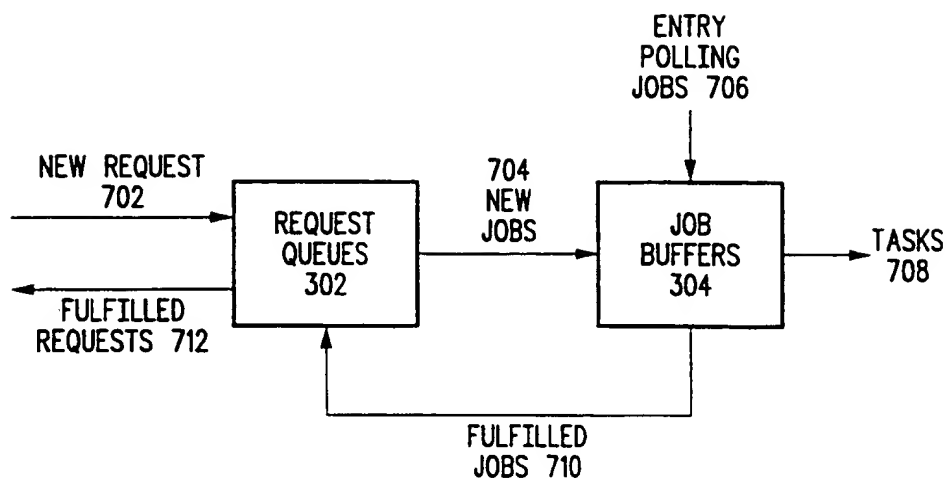
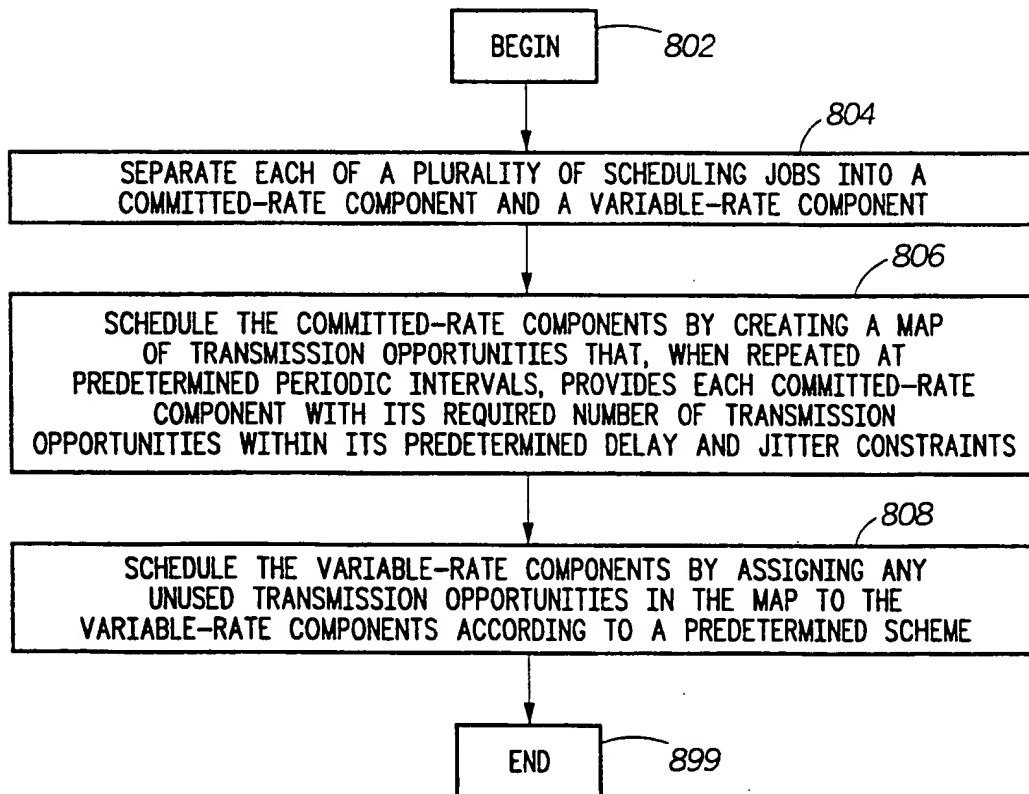
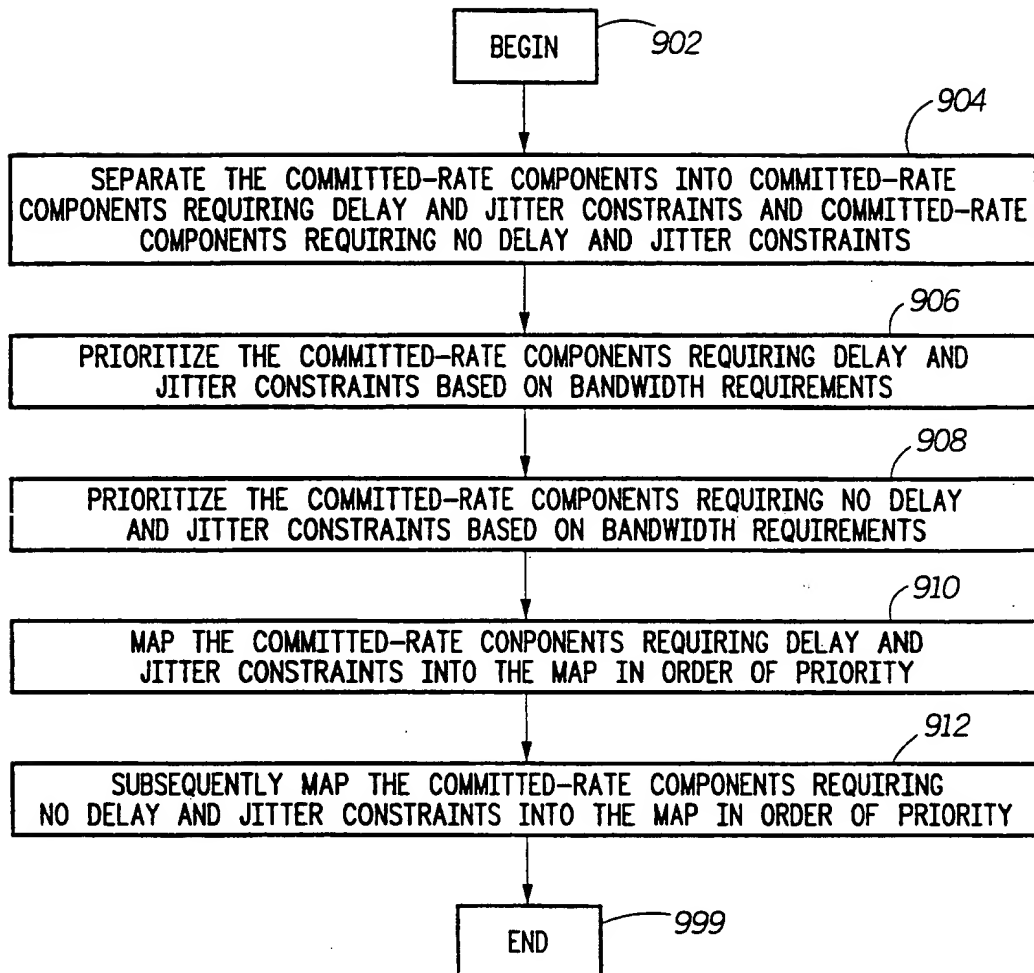


FIG. 6

6/7

*FIG. 7**FIG. 8*

7/7

*FIG. 9*

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US98/14893

A. CLASSIFICATION OF SUBJECT MATTER

IPC(6) :H04J 3/16

US CL :370/353, 468

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 370/352, 353, 412, 413, 414, 415, 416, 417, 418, 431, 465, 468

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 5,392,280 A (ZHENG) 21 February 1995, figures 1-4, col. 5, line 36 to col. 6, line 62.	1, 4 & 6-10
A	US 5,390,184 A (MORRIS) 14 February 1995, abstract, figures 3 & 5-6.	1, 4 & 6-10
A	US 5,008,878 A (AHMADI et al) 16 April 1991, figure 5, col. 8, lines 45-68.	1, 4 & 6-10



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents:	*T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
A document defining the general state of the art which is not considered to be of particular relevance	*X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
E earlier document published on or after the international filing date	*Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
L document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	*A* document member of the same patent family
O document referring to an oral disclosure, use, exhibition or other means	
P document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search

27 SEPTEMBER 1998

Date of mailing of the international search report

20 JAN 1999

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Facsimile No. (703) 305-3230

Authorized officer

HUY D. VU

Telephone No. (703) 308-6602

Joni Hill